

DRIVING ON THE SURFACE OF MARS USING THE ROVER CONTROL WORKSTATION

*Brian K. Cooper
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91109
818-354-6298
brian.k.cooper@jpl.nasa.gov*

ABSTRACT

On July 4, 1997 the Mars Pathfinder Spacecraft and Sojourner Microrover safely landed on the surface of Mars and began an ambitious mission of exploration. The Rover Control Workstation (RCW) was the software/hardware system used to control Sojourner. It provided the only means of creating the daily sequences of commands used to control all rover activities. Upon completion of the extended mission (30 days following landing), Sojourner had accomplished all project mission goals and continued to operate beyond expectations. It performed numerous technology experiments and proved the concepts for future planetary robotic missions. It has also provided large science returns by accurately placing its primary instrument, the alpha proton x-ray spectrometer on numerous Martian rocks and soil, to determine their elemental compositions as well as taking many high resolution rover-eye-view images of the surface using three on-board CCD cameras.

Using the RCW system to control Sojourner's movement across the terrain, the Rover Uplink Operations team led it to specific rocks and science targets. The RCW provided this team, consisting of the rover driver and sequence planner, the tools needed to create complex sequences of rover commands to accomplish the frequently changing daily science and technology goals. For mobility and articulation planning, it provided a unique interface consisting of mosaiced stereo windows displaying the panorama of Mars using camera images from both the lander and Sojourner. The rover driver used liquid crystal shuttered goggles to perceive stereo depth and a special six degree of freedom input device to move a stereo rover cursor on the screen. The system rendered this rover "CAD" model cursor in real time over the stereo image background, correctly simulating rover perspective, size, and appearance. The Uplink team used the RCW to make decisions about where to safely send Sojourner and what to do when reaching the goal. The RCW also provided a "virtual reality" type flying camera view of the surface using computer generated terrain models. This paper first provides an overview of the Rover mission. It then describes the design features of the RCW system and how it was used on a daily basis by the rover driver. The author is the creator of the RCW system and was the primary rover driver for Mars Pathfinder.

MISSION ELEMENTS

DESCRIPTION OF THE ROVER

The Sojourner Rover is a six-wheeled vehicle that is 68 cm long by 48 cm wide by 28 cm high, weighs 11 kg and contains a rocker-bogie chassis. The Warm Electronics Box (WEB) houses the on-board electronics, UHF modem, and batteries. A Gallium Arsenide solar panel and Material Adherence Experiment (MAE) are mounted on the top of the WEB. A pop-up UHF antenna, the Alpha Proton X-Ray Spectrometer (APXS), two black and white cameras, one color camera, and laser stripers (used for obstacle avoidance) are mounted outside the WEB. More detailed descriptions of the hardware can be obtained from [1].

MISSION SUCCESS CRITERIA

The Sojourner Rover had a combined scientific and technology mission. Once deployed on the surface of Mars, the success of the Rover was dependent upon completing a series of objectives: (1) collect data from one complete set of technology experiments including soil mechanics, material adherence, and wheel abrasion, (2) collect data from one APXS experiment, and (3) take an image of the Pathfinder lander to assess its post-landing condition. Sojourner completed all of these objectives, which constituted 90% mission success, during its nominal 7-Day mission. Sojourner continued to perform technology and science experiments as well as provide spectacular imagery throughout the remainder of the mission, satisfying the remaining 10% success criteria.

DESIGN CRITERIA

The goals of the mission and the operating conditions on the surface of Mars dictated the design goals for the RCW. Environmental issues included round-trip light time delays (approx. 22 minutes for a signal to travel to Mars and back at the start of the mission), limited communications opportunities (due to geometry and power considerations), visual aspects of the Martian environment such as lighting and atmospheric dust, and power management. Uncertainty in surface temperatures, local topography, the complexity of the landing site, and the orientation of the Pathfinder lander following the Entry-Descent-Landing sequence meant that we had to be prepared for the unexpected.

In order to achieve the mission success criteria the RCW design needed to include specific capabilities to support two primary functions: mobility planning and sequence generation.

MOBILITY PLANNING

Mobility planning was performed by the Rover Driver, and required a visual assessment of the Martian environment. Stereo vision was crucial to being able to detect hazards, evaluate the ability of the rover to fit through a given space, understand the risks associated with different paths, and determine the best approach attitude for reaching science targets. In addition to viewing the terrain in 3-dimensions (3-D), the Rover Driver also needed to interact in 3-D in order to designate the path for the Rover. Therefore, one major design requirement was the ability to accurately place a 3-D cursor in the position corresponding to the desired location on the Martian surface. This in turn led to the need to be able to provide ranging information based on stereo processing of images from the Pathfinder lander camera.

Through the RCW, the Rover Driver was able to designate a path for the Rover. The path consisted of a combination of low level move and turn commands and high level “go to waypoint” commands. In order to determine the best path, the driver needed to examine alternatives and evaluate “what-if” scenarios. To support this, the RCW needed to capture multiple paths and partial paths. When the best path was chosen, it was then integrated into the Rover command sequence, during the sequence generation activities.

COMMAND SEQUENCES

The operating conditions for the mission placed many constraints on Rover operations. Because there would be limited communications opportunities between the Pathfinder lander on Mars and the Earth-based teams, an entire day’s worth of activities needed to be planned in detail and sent to the Rover through the Pathfinder lander during each day’s single uplink opportunity. Near the end of each Martian day, Pathfinder sent back to Earth the results of that day’s activities. The Rover team would take these results and use them as the starting point for the next day’s activities.

The sequences sent to the Rover contained mobility and science instructions, as well as housekeeping instructions such as battery maintenance, heating instructions, and health checks. These commands needed to be integrated into a single sequence that took into account the time-of-day, sun angle, constraints due to competition for resources on the Pathfinder lander, and bandwidth limitations for sending data back to Earth. Part of sequence generation also included higher level mobility commands which set the level of automated safeguards that the Rover would use during a given traversal.

The average sequence would consist of approximately 200 commands, so the ability to edit, splice, and re-use portions of the sequence was important in streamlining operations. Each command could have a series of required parameters, optional parameters, and repeating parameters. Error checking to ensure that wrong values or data types weren’t specified was needed to reduce the potential for human error. The RCW enabled the human operators to work with English-language commands. These commands were then translated into binary codes that were transmitted to the Rover. As a safeguard measure, the binary codes were re-translated back to English by the RCW and an additional layer of error checking performed.

The available turn-around time between receiving the previous day’s data and sending the current day’s set of commands was very limited. The RCW user interface had to minimize operator workload, while enabling collaboration between operations team members and minimizing opportunities for error.

HARDWARE ARCHITECTURE

The RCW hardware configuration consists of two redundant collections of computer workstations and associated input/output devices. We could not afford any down time in case of hardware failure during the mission and so maintained two complete copies of the hardware and software. The software was designed and tested on a Silicon Graphics Crimson Reality Engine workstation. During the late stages of testing and during the actual mission we were fortunate enough to be able to borrow two of the newer Silicon Graphics Onyx2 Infinite Reality Engine models which gave RCW a significant boost in performance. These workstations each had one

gigabyte of RAM, 18 gigabytes of disk storage, and 64 megabytes of texture RAM for fast texture mapping. Input devices consisted of keyboard, mouse and a special 6-degree of freedom input joystick call Spaceball from Spacotec IMC Corp. Output consisted of 24 in. monitors and liquid crystal shuttered goggles from StereoGraphics Corp.

SOFTWARE ARCHITECTURE

The software consisted of over 40 thousand lines of C++ code with another 40 thousand lines of code using the X window Motif libraries. The Tools.h++ and Math.h++ class libraries from RogueWave software were used for string manipulation, linked list data structures and some math routines. RCW provided the user with simple to use graphical user interfaces (GUI) for each possible rover command. The majority of these GUI's were executed using the Motif library using it's windowed sliders, buttons, and text entry widget. Three dimensional (3D) visualization was done using the Open Inventor and OpenGL libraries. The stereo-in-a-window capability of the workstations allowed stereo images and stereo 3D rover model overlays to be viewed on the screen with goggles allowing depth perception by the user. This special view of the surface of Mars allowed the rover driver to detect hazards in the terrain and command the rover to stay clear of them.

STEREO PROCESSING

Before launch all cameras on the lander and rover were calibrated using a special procedure developed at JPL and accurate camera models were created from this process (see reference [2]). During operations, image products were generated by the Ground Data System and Multimission Image Processing System at JPL by transforming the raw telemetry packets into validated, linearized images. The RCW used these stereo images, along with the camera models, to create two mobility views: one for the "flying camera" and one to support Computer-Aided Remote Driving (CARD). In the CARD mode, images were displayed in a panoramic mosaic, which was viewed through stereo goggles to achieve the 3-D effect. Due to overlaps in the images, the RCW provided special adjustments that allowed the panels in the mosaic to be adjusted to minimize the visual breaks. Within this display, the operator moved at the Spaceball to control a scaled representation of the rover on the screen and through the terrain. This special cursor simulated the size and orientation of the Rover on the surface and provided additional visual cues as to distance and scale, which aided in path designation.

To drive the Rover, the driver used the Spaceball to move the virtual Rover through the terrain. When the operator wanted to tell the Rover to go to a specific location, he placed the icon of the Rover there and designated a "waypoint." In the user interface, these waypoints were indicated by "lawn-dart" icons. When the rover received these waypoints it would attempt to travel from one to another while using its on board sensors for navigation and obstacle avoidance. This provided a measure of safety in case the rover sensed, using its onboard hazard detection system, dangers unforeseen by the driver; the rover had the final say in this mode of driving and could choose to alter its path around large rocks or depressions and even give up and call home if certain programmable thresholds or a timeout was exceeded..

The second stereo interface was through the "flying camera." In this view, the stereo images were processed to create a 3-D terrain model. Topographic features such as the height of rocks and undulations in the surface were calculated and presented as either a plain x-y-z grid or as a

texture map on top of that grid. Because the view of the terrain was limited by the fixed location of the Pathfinder cameras, there were significant gaps in the data due to not being able to see behind rocks or other obstacles. Also, the resolution decreased as the distance from the cameras increased. Despite these limitations, however, the flying camera view proved to be extremely valuable in evaluating hazards.

For example, the RCW (in flying camera mode) played a key role in the assessment of hazards in deploying the Rover exit ramps. Throughout the cruise phase and the entry-descent-landing sequence, the Rover was stowed on one of the Pathfinder lander petals. In order to begin its surface operations, the Rover needed to exit the lander via one of two ramps, which would be deployed on the surface. The RCW was used to determine if it was safe to deploy the ramps. Based on landing conditions, it turned out that the airbags, which cushioned the Pathfinder landing, didn't retract completely and were actually located over the solar panels. This presented a hazard for Rover deployment (as well as other problems for the Pathfinder lander) so once the RCW system was used to ascertain that the airbags were in the way, a special operation was performed to retract the air bags. A special 3-D model of the lander was overlaid on the terrain model to enable this analysis.

Once the airbags were taken care of, the ramps were deployed, and once again the RCW flying camera view was used to assess the safety of using either ramp. Based on this view, it was clear that the forward ramp was actually sticking straight out and didn't reach the ground. The decision was made to exit via the rear ramp, which was considered significantly less risky. On July 5, 1997, the Rover safely descended to the surface of Mars to begin its explorations.

SEQUENCE GENERATION

Every action planned for the Rover on Mars was carefully choreographed and scripted using the RCW. The sequence generation requirements described previously were met by the RCW software and provided the capabilities necessary to support successful operations. The English-language command editing capability and graphical user interface proved to be a valuable tool for supporting the rapid turn-around required by the operations team and enabled them to work collaboratively for sequence development. The error checking, implemented using simple limit-checking, type-validation, and conditional techniques, provided a safeguard against input errors.

The goal of the sequence generation activity was to translate the science and technology goals into a set of commands that would balance these goals with the engineering risk to the Rover. Being able to play out what-if scenarios and explore options with respect to mobility and sequencing enabled the operations team to better meet the needs of the scientists and technologists. The general operations approach enabled by the RCW was based on years of experience with the operation of robotic vehicles [3, 4, and 5]. The specific capabilities of this generation of RCW were tested extensively under laboratory conditions as part of the development effort.

THE RCW DURING OPERATIONS TESTS

The Rover operations team, and all of their supporting tools, went through a rigorous testing program during the 7 months that the spacecraft traveled from Earth to Mars. As part of a comprehensive test program the Pathfinder project conducted a series of Operational Readiness

Tests (ORTs) which simulated critical phases of the mission and served as a dress rehearsal for the actual landing on Mars. In addition, the Rover team conducted a series of Rover Operational Readiness Tests (RORTs) to evaluate specific rover issues. During these tests, a “sandbox” which simulated Martian soil, rocks, and lighting was used to create a realistic environment. Full scale models of the Pathfinder lander and rover were placed in the sandbox where the operations team could not directly observe them. The only information available to the team was that which would be available during actual Mars operations.

The ORTs and RORTs served as true learning experiences for the Rover team and provided additional feedback on the design features of the RCW. As anticipated, the time pressures facing the operations team were intense. The time-saving features such as the GUIs for command generation and sequence editing and the built-in error checking helped the team to meet operations deadlines.

During the ORTs, the operations teams worked on “Martian” time. The Martian day (or “sol”) is 37 minutes longer than an Earth day. Therefore, each day of operations would be shifted later. Once the ORT began, the operations teams were working on a schedule defined by the Martian sol, around the clock, as if it were the real thing. The timeline was strictly followed and if the team couldn’t do their job in time to meet uplink or downlink deadlines – it didn’t get done and there was a one-day delay until the next communications opportunity.

For the Rover team, the ORTs represented the first opportunity to work with the scientists and begin translating their requests into workable sequences. While laboratory experiments and tests showed that the Rover could physically do the tasks necessary to conduct the experiments, the ORTs provided important insights into what made these tasks more (or less) valuable to the scientists. For example, during the ORTs, when the experiments were taking place under simulated Martian lighting conditions, it became clear that the Wheel Abrasion Experiment (WAE) had specific lighting requirements. These requirements implied that the Rover needed to not just be at the correct location, but also had to be pointed in the right direction to ensure proper lighting.

Another aspect of the ORTs that tested both the team and their tools was the project “Gremlin.” This role was played by a spacecraft engineer who, in setting up the sandbox configuration for a specific ORT, would introduce problems that the operations team would need to handle if they occurred for real on the surface of Mars. These problems included airbags covering a solar panel (which actually did occur), the lander being tilted at an extreme angle (which caused calibration difficulties), and the exit ramps for the Rover being blocked by hazards.

The demands of meeting the potential surface problems led to additional capability being added to the RCW. For example, special software was written late in the project to make sure that the team could easily evaluate the steepness of the ramp angles. This capability was put to use during the actual deployment of the ramps and in the decision of which ramp would be used to exit the lander. While the basic functionality for this existed in the RCW prior to the tests, the combination of the time constraints coupled with the likelihood of actually needing the capability made it worth the effort to create a fast and easy way of performing the function.

In general, the ORTs and RORTs enabled the operations teams to identify what functions needed to be done faster, more often, with higher fidelity, or with special inputs from the scientists. They also provided insights into how different aspects of the physical environment would affect their current capabilities. ORTs also provided valuable training time in learning how to interpret

the stereo imagery and to help the operators get used to working with the 3-D goggles. Finally, these tests showed the importance of having an operational environment that allowed the Rover team to work collaboratively in assessing the situation and making decisions. As an added benefit, many of the sequences and partial sequences developed during the ORTs were able to be reused during the actual mission.

ON THE SURFACE OF MARS

Despite all the attempts to simulate the Martian environment and its impact on operations, Mars itself and the public reaction to the mission held some surprises for the operations teams. With the Rover on the surface of Mars, the Rover team had to deal with communication outages, fluctuations in temperatures, power limitations, and degradation in performance of components. Planning activities needed to include contingencies and additional what-if scenario evaluation. Turn around times became even tighter due to the greater effort needed to assess the previous sol's activities.

The landing site on Mars presented significant variation in terrain, much more than was experienced in the sandbox tests or during field tests held in the California desert. In a relatively small area, there were major rock outcroppings, sand dunes, craters, and densely spaced collections of rocks. The landing site also provided great visibility and for the first time, the Rover operations team actually had to interpret stereo imagery that stretched to the horizon.

The variability of the terrain and the actual shapes and placement of rocks of scientific interest posed many challenges. The RCW flying camera and CARD modes were used extensively to determine the safety of paths through tightly constrained areas such as the "rock garden" and hazardous, complicated approaches to rock formations such as the "bookshelf". Arriving at the appropriate destination solved only part of the problem. The Rover had to also be in the correct orientation, whether it was to take a color image with the rear camera, stereo images with the front cameras, or (most difficult) place the APXS onto a rock face. The APXS presented a special challenge because of the height and surfaces of the rocks the scientists chose. If a rock was short and flat, or tall with a vertical face, the approach was relatively simple. However, a number of the target rocks were of a size where a topside or head-on placement of the APXS wouldn't work. The Rover team had to create a series of moves that presented the APXS at an appropriate angle while protecting the solar panels from being scratched, or accidentally driving up on the surface of the rock. One of the more-highly publicized misplacements occurred on the rock "Yogi" when, in an attempt to avoid a dangerous overhang, the Rover over-compensated and ended up driving on Yogi instead of placing the APXS on it.

During the early, critical operations, the intense public interest in the Pathfinder mission, and especially the Rover activities, added additional stress to produce certain data products with rapid turn-around times. One of the most popular products were the Rover Movies taken by the Imager for Mars Pathfinder (IMP) cameras. The "movies" were actually a series of time lapsed photos taken by the IMP cameras. Creating these movies required that the Rover team accurately predict the time and location of the rover during its moves. These predictions were then used by the lander science operations team to create special sequences for the IMP cameras that had to be synchronized with the Rover actions. The ranging features of the RCW were used to provide pointing information to the cameras. The sequence generation capability was used to identify when key events were supposed to occur and provide that information also.

The RCW was used throughout the entire Rover mission. Every command that was sent to the Rover was generated by the Rover operations team using the RCW. In addition, when not being used to support operations, the RCW was used for demonstrations to visiting dignitaries, to provide visual displays for the media, and to work with scientists on evaluating longer-range objectives. By Sol 83 when contact with the Pathfinder lander was lost, all the Rover mission objectives had been met, and the Rover had operated for over 10 times its required lifetime.

FUTURE DIRECTIONS

The success of the Mars Pathfinder mission and the Sojourner Rover are the standard by which future missions will be judged. Current activities are focusing on the development of much larger, more ambitious Rover missions for the Mars Surveyor program. The Rover Control Workstation (RCW) for these missions will be based on the Pathfinder RCW, but will incorporate many new features to support the more complicated rover and mission, as well as taking advantage of improvements in technology and the lessons learned on Pathfinder.

ACKNOWLEDGEMENTS

I would like to thank the following individuals for their help in making the RCW system a success: Andy Mishkin, Jack Morrison, Rick Welch, and Frank Hartman.

The work described in this publication was carried out by the Jet Propulsion Laboratory, California Institute of Technology under a contract with the National Aeronautics and Space Administration.

REFERENCES

- clear
0691-
Bul*
- [1] Mishkin, A. H., Morrison, J. C., Nguyen, T. T., Stone, H.W., Cooper, B.K., and Wilcox, B.H. (1998). "Experiences with Operations and Autonomy of the Mars Pathfinder Microrover." Proceedings of the 1998 IEEE Aerospace Conference, Mar. 21 - 28, 1998, Snowmass, CO.
 - [2] Y. Yakimovsky and R. Cunningham. A system for extracting three-dimensional measurements from a stereo pair of TV cameras. Computer Graphics and Image Processing, 7:195-210, 1978.
 - [3] Wilcox B., L. Matthies, D. Gennery, B. Cooper, T. Nguyen, T. Litwin, A. Mishkin, H. Stone, "Robotic Vehicles for Planetary Exploration", Jet Propulsion Laboratory, IEEE International Conference on Robotics and Automation, Nice, France, May 12-14, 1992.
 - [4] Cameron, J., B. Cooper, R. Salo, B. Wilcox, "Fusing Global Navigation with Computer-Aided Remote Driving of Robotic Vehicles", presented at the SPIE International Conference, Boston, MA, October 26-31, 1986.
 - [5] Wilcox, B., B. Cooper, R. Salo, "Computer Aided Remote Driving", presented at the Association for Unmanned Vehicle Systems (AUVS-86) Conference, Boston, MA, July 21-23, 1986.